International Journal of Forecasting 14 (1998) 313–322

The impact of information of unknown correctness on the judgmental forecasting process

William Remus[a,*], Marcus O'Connor[b], Kenneth Griggs[c]

[a]*University of Hawaii, 2404 Maile Way, Honolulu, HI 96822, USA*
[b]*University of New South Wales, Sydney, Australia*
[c]*MITRE, San Diego, USA*

Abstract

Modern organizations are awash in information. This information can be very useful and have a major impact on the organization and its future. Much of this information, however, is of unknown correctness. This study investigates whether such information can be effectively used to aid people in forecasting changing time series. In the first section of this study, correct information improved the subjects' forecast accuracy more than incorrect information or no information. Incorrect information, however, resulted in no worse forecast accuracy than no information. The subjects then continued making forecasts as the level of correctness changed (for example, from correct to incorrect information). Those subjects who received incorrect messages at any time during the experiment made less accurate forecasts near the turning point but equivalent forecasts in the long run.

Keywords: Judgmental forecasting; Adjusting forecasts; Information; Time series; Structural change

1. Introduction

Our world is characterized by change and uncertainty. This presents a considerable challenge to forecasters. When conditions are changing, it is essential to consider all information, including information of unknown correctness, as part of the forecasting process. Forecasters cannot rely simply on the re-occurrence of past patterns or trends.

Accurate information can be provided to aid the forecaster; for example, there are data services which provide data on the economy and there is much accurate internal data on many aspects of the organization. There is also information of unknown correctness that comes from the sales force, competitors, personal relationships with others inside the organization, trade publications, rumors, and the grapevine. This potentially relevant information is of varying specificity and unknown reliability and validity. What happens when forecasters try to use the latter information?

This article is about the impact of information of unknown correctness, such as the rumors discussed above, on the accuracy of judgmental forecasts particularly when turning points occur. Our approach is first to illustrate the use of information in forecasting by describing its use in the product forecasting task. Next we review the relevant literature. Finally, we describe the experimental design, report the

*Corresponding author. Tel.: (1-808) 956-7608; fax: (1-808) 956-9889; remus@hawaii.edu

0169-2070/98/$19.00
PII: S0169-2070(98)00002-8

results, and discuss the implications of an experiment about the impact of information of unknown correctness.

1.1. Information use in product forecasting meetings

Product forecasting requires the use of both information known to be correct and information of unknown correctness. Information known to be correct may include information such as past production and sales of a product. Information of unknown correctness would include information such as assessments of competitors' future production and promotional activities.

A common method of incorporating information of unknown correctness into the forecasting process is through meetings of product forecasting committees. Typically, these meetings are attended by marketing, sales, manufacturing, inventory and finance personnel. It is at these meetings that various points of view are expressed, new information is provided, interpretations are presented and a forecast made. For example, information about forthcoming advertising schedules, maintenance of machinery, and problems with distributors are vital pieces of information that need to be incorporated into the forecast (Lawrence et al., 1995).

The distinguishing characteristic of much of the information provided at the product forecasting meetings is that it is often subject to varying interpretations. Also, by nature, it is difficult to establish whether the information is correct or incorrect. Past time series characteristics are information known to be correct but are only considered as a base from which to incorporate additional information.

Information presented at product forecasting meetings can be critical in detecting changes in the environment. In practice, changes are largely detected or anticipated by using information of unknown correctness. Information that changes are occurring impacts not only the forecast but also the forecasting model to be selected (Collopy and Armstrong, 1992). In spite of the importance of detecting such changes, Collopy and Armstrong (1992) point out that little attention is given to developing formal methods to detect change.

Section 2 reviews the literature on judgmental and statistical forecasting when conditions change and the literature on the importance of information in organizations. This review suggests a laboratory experiment designed to measure the impact of information of unknown correctness on the forecasting of changing time series.

2. Background

Surveys of the practice of forecasting have repeatedly found human judgment to be the overwhelming choice for forecasting (especially sales forecasting, Dalrymple, 1987; Sanders and Manrodt, 1994; Taranto, 1989). This is true even after statistical approaches have been tried (Lawrence, 1983). Apparently, managers feel more comfortable dealing with their own or a colleague's estimates than with statistical models (Langer, 1975).

2.1. Information use in judgmental forecasting

The widespread use of human judgment to make business forecasts might be rationalized in two ways. The first is that people might be better able to detect changing patterns in the time series than are statistical forecasting methods. A recent experiment (O'Connor et al., 1993) examined people's ability to detect and respond to unexpected changes in noisy time series data. The subjects were given the raw time series data graphically displayed and were asked to forecast the data point by point. The subjects were unable to outperform simple time series forecasting models when unexpected time series turning points occurred. Moreover, the subjects displayed significant bias towards overreaction to random variation in time series. Thus, the first rationale lacks support.

A second rationale for the use of human judgment is that people might be more able to integrate outside (i.e. non-time series) information into the forecasting process (Johnson, 1988, p. 213). Often this information can indicate that the time series pattern is changing.

There is support for such an assertion in the forecasting literature; a number of studies have shown that non-time series information can contribute to forecast accuracy. Edmundson et al. (1988); Sanders and Ritzman (1992) found that contextual knowledge of the forecast variable was a significant contributor to accuracy in product forecasting. This was especially true for products that were important to the organization. Intimate knowledge of the underlying forces shaping the product (e.g. market plans and competitor issues) is crucial in deciding about future direction of the forecast variable. Reliable, non-contextual information can also improve the accuracy of forecasts. A recent experiment (Remus et al., 1995) found that more accurate forecasts resulted from increasingly reliable information. However, the same experiment also found that such integration of reliable information was not optimal.

The kind of information that is integrated into the forecasting process to improve forecasts is often informal information; that is, information that is not communicated through formal organizational channels. Organizations are awash in informal information; five out of six messages in an organization are informal ones and often travel through the grapevine (Lewis, 1980). Such messages carry more accurate than inaccurate information (Rudolph, 1973), but the messages are typically incomplete (Lewis, 1980). In a number of cases, these informal messages assume the characteristics of 'rumors' (Rosnow, 1992). The regular product forecasting meetings often deal with such informal communication and rumors. Such informal information could contribute significantly to improving forecast accuracy (Edmundson et al., 1988; Sanders and Ritzman, 1992).

As noted above, informal information is often of unknown correctness. In the product forecasting context, this is most obvious with the speculations about competitor actions and about the likely success of marketing efforts, especially advertising. Often the consequences of informal information are also unknown. For example, sometimes there are also disagreements in the product forecasting meeting on the likely impact of the issues raised on future sales. So the meeting needs somehow to factor into the final forecast some assessment of the two kinds of unknowns associated with the informal information. We will focus on the first kind of unknown; that is, information of unknown correctness.

Assessments of correctness are not made in temporal isolation; that is, the correctness of today's information is affected by the correctness of similar information in the past. For example, if sales estimates presented by marketing were too optimistic in the past, those at the product forecasting meeting will tend to discount marketing's current sales estimates. Thus, a complete research design should include not only the correctness of the information presented but also sequencing of the information.

In the research about to be described, we will first focus on whether judgmental forecasters are effective in using information of unknown correctness when forecasting changes in time series. In the second part of the experiment, we will examine if judgmental forecasters are affected by the correctness of previous information about changes in time series. The results of the research will be reported following a section describing our experimental design.

3. Research design

3.1. Time series

Time series containing discontinuities were generated for this experiment. As in our prior work (O'Connor et al., 1993; Remus et al., 1995), each time series was divided into three contiguous segments:

- Segment 0 (periods 1–20):
  This segment was 20 periods of historical data generated with a base of 100 and error added. This data was displayed to the subjects so they could assess the initial characteristics of the time series.
- Segment 1 (periods 21–28):
  This segment was a continuation of the series as displayed for the first 20 points (segment 0). However, now the subjects made forecasts in each of the periods from 21–28. In this way, they became accustomed to forecasting the series.
- Segment 2 (periods 29–38):
  In this segment, the discontinuity was introduced; the first departure from the segment 1 flat pattern occurred in period 29. The discontinuities were a ramp upward (growth) or a ramp downward

(decline). In addition, we included a series form (as a control) where the series did not change.

The randomness in the series was drawn from uniform distributions with a MAPE (mean absolute percentage error) of 5%; every subject received time series with the above underlying levels but differing patterns of randomly generated error. The MAPE used is representative of error rates in sales forecasting (Dalrymple, 1987; Taranto, 1989) and of our earlier experiments.

3.2. Information provided

To test the major research questions, information about the discontinuities in the time series was provided to the subjects. The information was presented in the form of messages displayed as the discontinuity was occurring; these messages were displayed for four periods (periods 28–31). The information was expressed as 'rumors' since the term rumors conveys that the messages are of unknown correctness. The messages are shown in Table 1.

Some information messages correctly told what discontinuity would take place; we call these correct information messages. Other information messages provided incorrect information about the discontinuities that might occur; we term these incorrect information messages. Lastly, in some cases we provided no information about the discontinuity as a control condition.

As in our earlier work (O'Connor et al., 1993; Remus et al., 1995), all subjects were presented with 'Up', 'Down' and 'Flat' time series. Up series had a discontinuity at the start of segment 2 as the flat series (mean 100) changed to a ramp upward series growing at 2 units per period. If each of these series were quarterly data, the ramp up would correspond to an 8% annual increase. For the first four periods of the discontinuity the time series data were accompanied by message A1, A2 or B1, thus creating correct information, incorrect information, or no information, respectively. Similarly the down series declined at 2 units per period for an overall 8% annual decrease and for four periods the time series data were accompanied by message A2, A1 or B1, creating correct information, incorrect information, or no information, respectively. The flat series neither increased nor declined and for four periods the time series data were accompanied by message A3, A4 or B1, creating correct information, incorrect information, or no information, respectively. All these series were masked with error as described earlier.

Since subjects' prior experience with the information provided could effect the quality of their subsequent forecasts, the order of presentation of the information might make a difference. We have chosen to control for this effect in our experimental design by segmenting each treatment into two parts: the first six series where the subjects received messages of the same level of correctness and the last six series where the messages were of a different level of correctness. For example, in the case of treatment T2B below, we were interested in whether the first six time series with incorrect messages would affect forecasting accuracy in the next six time series where there was no information.

The subjects were assigned to one of five treatments that defined the order in which the series were presented and the number of presentations of each information type. The order of presentation and the number of series in each treatment were as follows:

Table 1
Details of time series directions at the change point and information conditions

These are the messages presented to the subjects:

A1 Message: There is a rumor that the series will now begin growing at 2 units per period.

A2 Message: There is a rumor that the series will now begin declining at 2 units per period.

A3 Message: There is a rumor that the series will continue as it had in the last few periods.

A4 Message: There is a rumor that the series will no longer continue as it had in the last few periods; that is, there is a rumor that the series might either grow or decline.

This is the no information message:

B1 Message: We have no information on whether this series will grow, decline, or remain stable.

The combination of the message and the direction of the turning points establishes the correctness of the message. For example, using A1 with an Up series results in a correct message; using A2 with a Down series results in an incorrect message.

	Information presented at Point 28		
Pattern	Correct	Incorrect	No information
UP	A1	A2	B1
DOWN	A2	A1	B1
FLAT	A3	A4	B1

- T1: Six time series with Correct Information
  T1A: followed by six time series with Incorrect Information
  T1B: followed by six time series with No Information
- T2: Six time series with Incorrect Information
  T2A: followed by six time series with Correct Information
  T2B: followed by six time series with No Information
- T3: Twelve time series with No Information

Treatments T1B and T2B measured the persistence of the impact of message correctness when no meaningful additional information was presented. T1A and T2A measured the persistence of the impact of message correctness when the correctness of the messages changed. T3 was a control treatment for both of the above comparisons. Within each of the treatments, the initial six series had two up, two down, and two flat series; these were randomly presented. These six series were followed by six additional series with two up, two down, and two flat series again presented in random order.

We might note that the above procedures emphasize experimental control. The messages are simple and of unknown correctness. They are not ambiguous, complex informal messages but clearly stated messages identified as rumors. While this approach may sacrifice generalizability at the cost of control, it is more likely that such a controlled experiment will find interesting results that might be lost in the error variance of an experiment using ambiguous, complex messages.

3.3. Details of the data-gathering procedures

The display of data and gathering of the subjects' forecasts was done with Hypercard software using a mouse interface running on Macintosh computers. After viewing a graphical display of the first 20 points of the series, subjects were required to forecast the next point in the series. When making their forecasts, subjects needed only to use the mouse to 'point' their forecasted value on the graphical display and 'click' to record that value. Following the forecast for the current point, the actual value of the time series would be graphically displayed. The cursor then moved to the right ready for the next forecast. This sequence continued until the end of the series. The series were presented in a unique random order to each of the subjects.

The 49 subjects for the experiment were undergraduate students at the University of Hawaii; there were 10, 9, 11, 10 and 9 subjects, respectively, in the five treatments. The subjects were recruited from an Operations Research course which covered time series forecasting in the two weeks prior to the experiment. In addition, these subjects were trained in the use of the software prior to beginning the experiment.

The 49 subjects were required to forecast all series after they had gained practice with a sample series (not included in the analysis). The subjects were told that some series contained a discontinuity and that their task was to be on the look-out for such discontinuities. Further, they were told that information about future discontinuities would periodically be displayed on their computer screen; they were told that some of these information messages were correct and others incorrect. The average time taken to complete all the tasks was about 1 h.

Subjects received both prizes and course credit for participating in the experiment. A prize was \$20 given for the best forecasting performance in each of the five treatments. The winner of each prize was the subject who had the lowest mean absolute percentage error (MAPE) in the treatment. All subjects earned full course credit for their participation.

4. Results

As described in Section 3, the subjects forecasted sequential points in a time series; thus, the forecasts were repeated measures of their forecasting skills. To address this, we analyzed the data using Repeated Measures Analysis of Variance. Our measure of forecasting accuracy was the mean absolute error of the forecasts (using MAPE yields the same general results).

4.1. Does the correctness of the initial information have an impact?

First, we examined the impact of the correctness of the initial messages on forecasting accuracy. To do that we compared T1, T2 and T3 following the

Table 2
Mean absolute percentage error across treatments (T1, T2, T3) and time series direction during the first six time series using Repeated Measures Analysis

Tests of significance for absolute forecast error					
Source of variation	Sum squares	*DF*	Mean square	*F*	Significance
Tests of between-subjects effects					
Treatment	518.00	2	259.00	6.26	<0.002
Direction	671.71	2	335.85	8.12	<0.001
Treatment by direction	48.87	4	12.22	0.30	<0.881
Error	10 136.19	245	41.37		
Within-subject effects					
Time	1404.28	10	140.43	4.59	<0.001
Treatment by time	832.59	20	41.63	1.36	<0.130
Direction by time	964.39	20	48.22	1.58	<0.049
Treatment by direction by time	1445.89	40	36.15	1.18	<0.202
Error	74 914.21	2450	30.58		

turning point in the first six time series. As shown in Table 2, both Treatment (T1 Correct Information, T2 Incorrect Information and T3 No Information, $p<0.002$) and the time series Direction ($p<0.001$) were significant; their interaction was not significant ($p<0.881$). The treatment means are shown in Fig. 1.

The method of multiple comparisons (Pedhazur and Schmelkin, 1991, pp. 482–490) was used to explore the relationships in the data. Correct Information (absolute error of 7.32) led to significantly better forecasting accuracy than Incorrect Information (absolute error of 8.24) ($t=-3.52$, $p<0.001$). Correct Information (absolute error of 7.32) also led to significantly better forecasting accuracy than No Information (absolute error of 7.89) if one makes the one tail hypothesis that Correct Information will lead to more accurate forecasts ($t=-1.720$, $p<0.036$). Incorrect Information and No Information were not significantly different ($t=0.931$, $p<0.352$). Also accuracy was significantly better when predicting flat series than when predicting upward ($t=-3.301$, $p<0.002$) or downward ($t=-3.656$, $p<0.001$) trending series. There was no significant difference when predicting downward and upward series ($t=0.517$, $p<0.606$).

There were also two significant repeated measurement effects shown in Table 2. An inspection of the diagnostic information provided by the Repeated Measures procedure showed the effects were largely due to the difference in the level of forecast error between periods 28 and 31, when the message was displayed, and periods 32 to 38 by which time the subjects had seen data confirming change in direction. This led to the significant Time effect ($p<0.001$). The Direction by Time ($p<0.049$) interaction resulted from differing behavior when flat series occurred than when the upward or downward discontinuities occurred.

4.2. Did the information provided impact turning point forecast accuracy?

Next, we compared the treatments (T1A, T1B, T2A, T2B and T3) during the second six series; that is, after the level of correctness changed. In treatment T1A, after six time series with correct messages, there were six time series with messages providing incorrect information. In treatment T1B, after six time series with correct messages, there were six time series providing no information. In treatment T2A, after six time series with incorrect messages, there were six time series with messages providing correct information. In treatment T2B, after six time series with incorrect messages, there were six time series with no information. T3 subjects never were provided with information of changes in the time series. As shown in Table 3, we found no significant effects among the treatments ($p<0.982$) although Direction was again significant ($p<0.001$). There was no significant interaction of the variables ($p<0.877$).

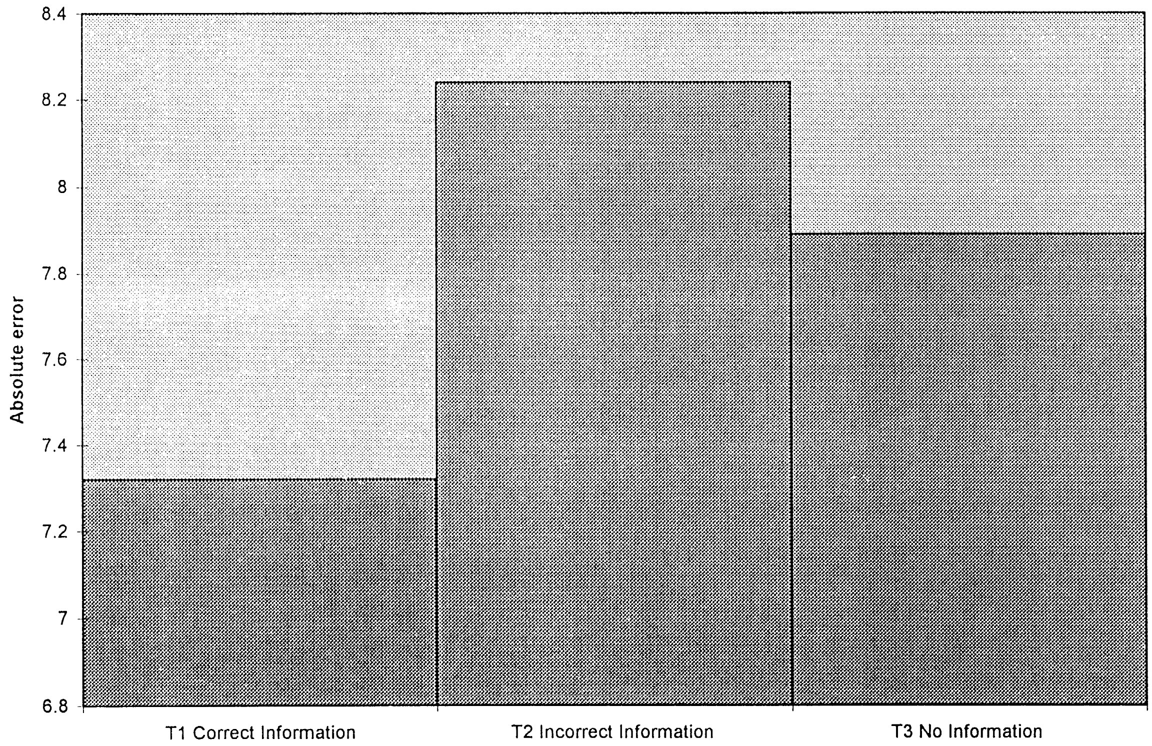

Fig. 1. Mean for each treatment during the first six series.

Table 3
Mean absolute percentage error across treatments (T1A, T1B, T2A, T2B and T3) and time series direction during the second six time series using Repeated Measures Analysis

Tests of significance for absolute forecast error					
Source of variation	Sum squares	*DF*	Mean square	*F*	Significance
Tests of between-subjects effects					
Treatment	14.44	4	3.61	0.10	<0.982
Direction	1523.63	2	761.81	21.55	<0.001
Treatment by direction	132.88	8	16.61	0.47	<0.877
Error	9084.96	257	35.35		
Within-subject effects					
Time	1135.25	10	113.52	3.73	<0.001
Treatment by time	1095.13	40	27.38	0.90	<0.650
Direction by time	1099.76	20	54.99	1.81	<0.015
Treatment by direction by time	2549.48	80	31.87	1.05	<0.365
Error	78 149.76	2570	30.41		

In multiple comparisons based on treatment, no contrasts were significant. The multiple comparison analysis for Direction was similar to that in the first six series. The accuracy was significantly better when predicting flat series than when predicting upward ($t=2.04$, $p<0.042$) or downward ($t=-4.418$, $p<0.001$) trending series. Also subjects had significantly more accurate forecasts when predicting upward series than when predicting downward series ($t=6.427$, $p<0.001$).

There were significant repeated measurement effects again. An inspection of the diagnostic information again showed the effects were largely due to the difference in the level of forecast error between periods 28 and 31, when the message was displayed, and periods 32 to 38 by which time the subjects had seen data confirming change in direction. This led to the significant Time effect ($p < 0.001$); the Time by Direction ($p < 0.015$) interaction was again significant.

4.3. Did the incorrect information have short-term effects on forecast accuracy?

Fig. 2 shows the absolute error immediately after the turning point (periods 28–31). Those treatments in which incorrect information had occurred at any time (T1A, T2A and T2B) appeared to have greater forecasting error than did the treatments in which no incorrect information had been given (T1B and T3). To explore this, we collapsed the five treatments into the two groups described above and performed the statistical test for periods 28–31 again. In this analysis the groups that had received incorrect information at any time (absolute error of 7.20) were significantly less accurate ($p < 0.003$) than the groups that had at no time received incorrect information (absolute error of 6.52). Direction was again significant ($p < 0.010$) but the interaction of Direction and Treatment was not ($p < 0.251$). All tests involving time were not significant at the 0.05 level.

Thus, there is some evidence that having previously received incorrect information reduced forecast accuracy at the turning point even though overall the groups did not differ. Thus, the impact of incorrect messages rapidly disappeared.

5. Discussion

The correctness of the information provided to the subjects during the first six series had an impact on forecasting accuracy as measured by absolute error.

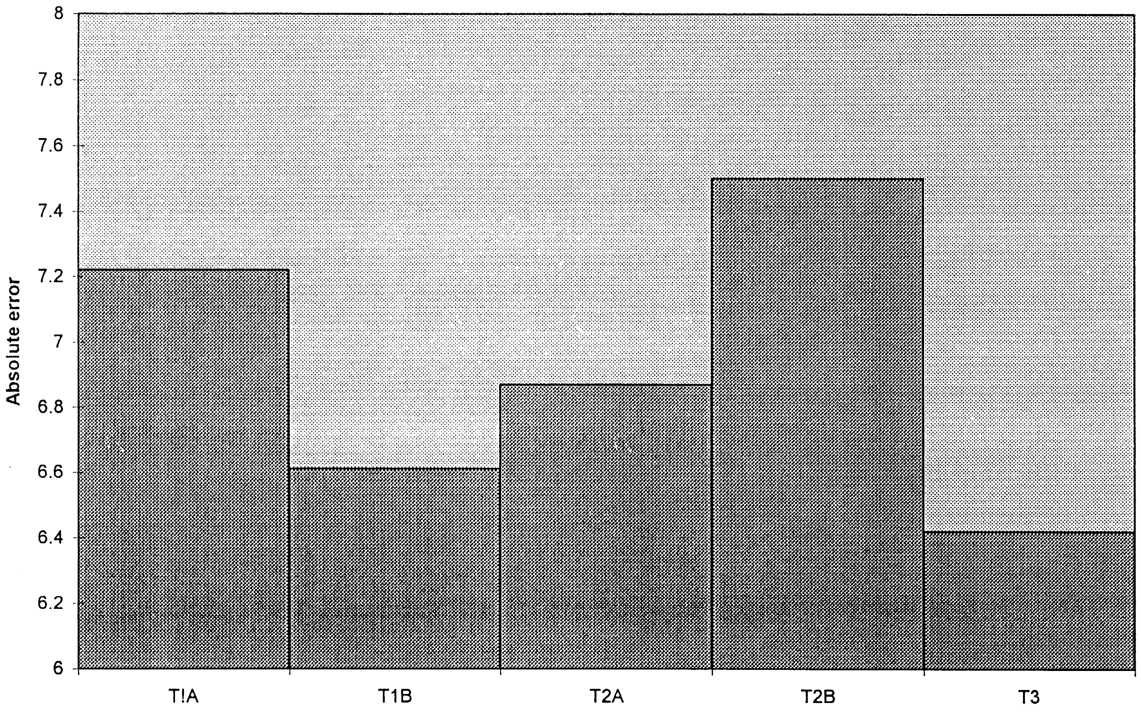

Fig. 2. Means for treatments immediately following the turning point in the second six series.

Correct information led to better forecasting accuracy than incorrect information or no information. Unexpectedly, the incorrect information led to the same forecast accuracy as no information.

The subjects effectively used the correct messages to improve their forecasts. The lack of difference between incorrect and no information seemed to result from the way that the subjects processed the messages that were incorrect. The subjects that received only correct messages unhesitatingly used the correct information. The subjects that had previously received an incorrect message looked for confirming information before acting on the message. In this experiment, the message correctness could be assessed by making several conservative forecasts and observing the actual data in the next several periods. Handling incorrect messages in this way gave results as good as having no information.

We next introduced a change in the correctness of the messages and examined the subjects' reactions. Overall, there were no significant differences among the treatments. Even correct information during this second set of six series did not lead to improved forecast accuracy. However, those treatments where subjects had at some time received incorrect information appeared to have more forecast error than those treatments where subjects had received no incorrect information at all immediately after the turning point. The effect quickly disappeared.

The short-term effect described above was important in two ways. First, this effect suggests that, when people have previously received incorrect information, they are conservative in responding to new information even if correct. Thus, incorrect information does carry over but rapidly disappears. This differs from the first part of the experiment, where one treatment only got correct messages and so the subjects responded to them without inhibition.

Second, the existence of the short-term effect assures us that the information conditions had an impact on the second part of the experiment. Thus, the overall finding of no significant difference between the five treatments cannot be attributed to an ineffective experimental manipulation.

The subjects had particular difficulty with downward sloping series. This finding is not unusual and is replicated in a number of studies (e.g. Lawrence and Makridakis, 1989; Winton and Edmundson, 1992; O'Connor et al., 1993). The reasons for this effect are examined in some detail in O'Connor et al. (1997).

Several groups of researchers (Edmundson et al., 1988; Sanders and Ritzman, 1992) have noted improved accuracy when contextual information is provided. Implicit in other researchers' work is the assumption that the contextual information was correct. Our work, however, suggests that the contextual knowledge must be known to be correct to use fully the information to improve forecast accuracy.

There are two important caveats about our experiment that relate to the way in which we operationalized information of unknown correctness and how we presented the sequence of series. The information was in the form of a context free 'rumor'. We did this to provide a clear message about the change. Had we put the clear message in a context, the subjects may have tried to use the contextual information to assess the correctness of the message (e.g. source credibility). Our approach, while not especially realistic, avoided ambiguity and misinterpretation of the message.

We operationalized information of unknown correctness as a series of six messages of one level of correctness followed by six messages of another level of correctness. In the real world the accuracy of messages is not so neatly ordered. We did this so that, if there was a carryover effect, we would find it. We found only a short-term carryover of the level of correctness even with the six messages in a row of the same correctness level. Given the relative artificiality of initially being given six messages in a row of the same correctness, this small carryover effect may be even smaller in the real world.

6. Conclusions

The results show that people can effectively use information of unknown correctness to improve their forecast accuracy. Generally, people can use that information when correct and they also are not necessarily misled by incorrect information. Putting that another way, incorrect information often leads to no worse results than having no information. There is also a carryover effect of incorrect information that is only short-term and near the turning point. The incorrect information is rapidly discounted so that no long-term loss of forecasting accuracy results.

References

Collopy, F., Armstrong, J.S., 1992. Expert opinions about extrapolation and the mystery of the overlooked discontinuities. *International Journal of Forecasting 8*, 575–582.

Dalrymple, D.J., 1987. Sales forecasting practices: Results of a United States survey. *International Journal of Forecasting 3*, 379–392.

Edmundson, R., Lawrence, M.J., O'Connor, M.J., 1988. The use of non-time series information in sales forecasting: A case study. *Journal of Forecasting 7*, 201–211.

Johnson, E.J., 1988. Expertise and decision under uncertainty: Performance and process. In: Chi, M.T.H., Glaser, R., Farr, M.J. (Eds.), *The Nature of Expertise*, Lawrence Erlbaum, Hillsdale, NJ, pp. 209–228.

Langer, E.J., 1975. The illusion of control. *Journal of Personality and Social Psychology 32*, 311–328.

Lawrence, M.J., 1983. An exploration of some practical issues in the use of quantitative forecasting models. *Journal of Forecasting 1*, 169–179.

Lawrence, M.J., Makridakis, S., 1989. Factors affecting judgmental forecasts and confidence intervals. *Organizational Behavior and Human Decision Processes 43*, 172–187.

Lawrence, M.J., Edmundson, R.H., O'Connor, M.J., 1995. A field study of sales forecasting: Its accuracy, bias, and efficiency, *Working Paper*, University of New South Wales.

Lewis, P.V., 1980. *Organizational Communication*: *The Essence of Effective Management*, Grid, Columbus, OH.

O'Connor, M.J., Remus, W.E., Griggs, K., 1993. Judgmental forecasting in times of change. *International Journal of Forecasting 9*, 163–172.

O'Connor, M.J., Remus, W.E., Griggs, K., 1997. Going up–going down: How good are people at forecasting trends and changes in trends? *Journal of Forecasting 16*, 165–176.

Pedhazur, E.J., Schmelkin, L.P., 1991. *Measurement and Analysis*: *An Integrated Approach*, Student Ed., Erlbaum, Hillsdale, NJ.

Remus, W.E., O'Connor, M.J., Griggs, K., 1995. Will reliable information improve the judgmental forecasting process?. *International Journal of Forecasting 11*, 285–293.

Rosnow, R.L., 1992. Rumor as communication: A contextualist approach. In: Hutchinson, K.L. (Ed.), *Readings in Organizational Communication*, W.C. Brown, Dubuque, IA.

Rudolph, E.E., 1973. Informal human communication systems in a large organization. *Journal of Applied Communications Research 1*, 7–23.

Sanders, N., Manrodt, K., 1994. Forecasting practices in US corporations: Survey results. *Interfaces 24*, 92–110.

Sanders, N., Ritzman, L., 1992. The need for contextual and technical knowledge in judgmental forecasting. *Journal of Behavioral Decision Making 5*, 39–52.

Taranto, G.M., 1989. *Sales forecasting practices: Results from an Australian survey*, Unpubl. thesis, University of New South Wales.

Winton, E., Edmundson, R., 1992. An examination of the judgmental identification and extrapolation of trend in time series, *Working Paper*, University of New South Wales.

Biographies: William REMUS is a Professor in the College of Business Administration at the University of Hawaii. His research interests focus on the process of human judgement as used in forecasting and decision making.

Marcus O'CONNOR is an Associate Professor in the School of Information Systems at the University of New South Wales, Sydney, Australia. His research interests center on the process of forecasting, especially the way judgement is incorporated into the final forecasts.

Ken GRIGGS is researcher at MITRE in San Diego. His research interests focus on object-oriented languages and systems analysis.